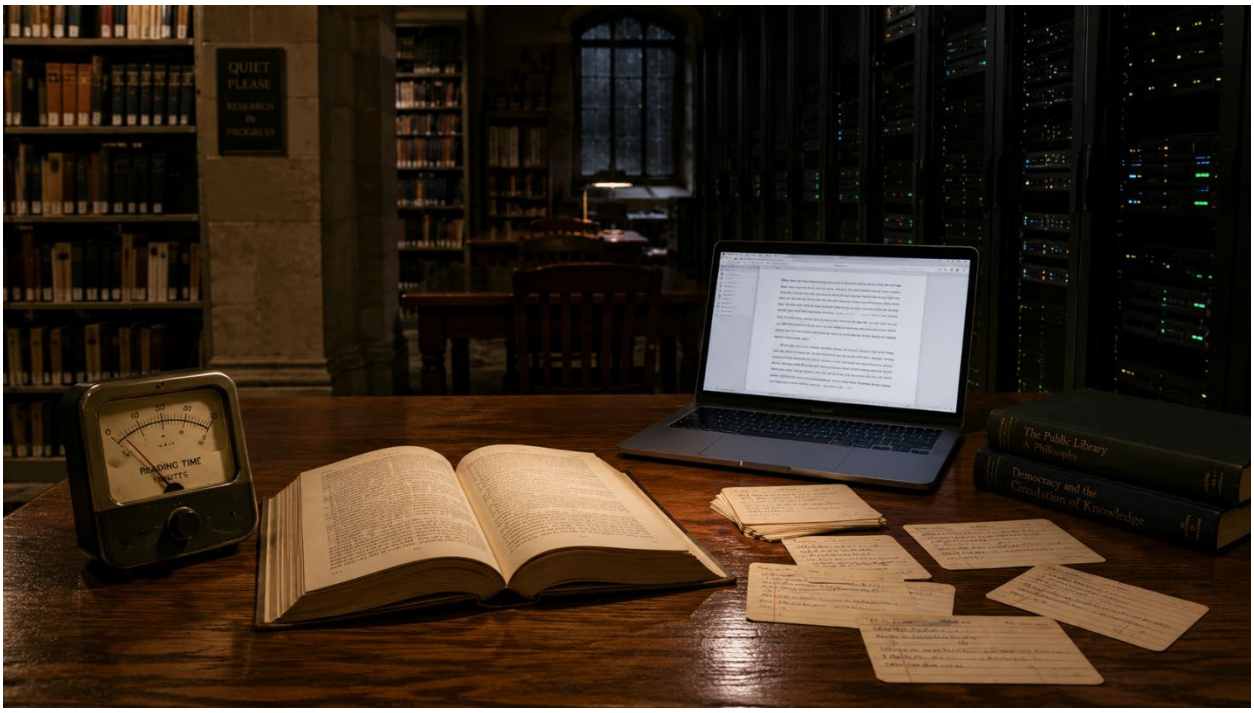

The Next Paywall Will Think: AI, Access, and the Brief Democracy of Machine Attention

KC Clarke

My vantage point is strange enough to be useful. I have spent much of my adult life inside systems that do not usually touch: poetry, arts administration, military memory, criminal reentry, nonprofit fundraising, humane law enforcement, museums, libraries, public festivals, donor rooms, board rooms, bookmaking, and court-adjacent work. I have been close enough to money to know how it speaks, close enough to art to know how it resists, close enough to public service to watch language become procedure, and close enough to institutional ambition to recognize when a promise is already hardening into a gate.

I do not have a crystal ball. I have not found a shop that sells one that works. What I have is a filter. It was built over years of watching institutions say one thing in public while another set of pressures moved underneath. Once, while interviewing a senior military leader on camera, I heard him say he could see war flags on the horizon. He was not claiming prophecy. He had simply seen enough to recognize the weather. That is the kind of knowledge I mean here.

When I look at AI, I do not see only a writing tool, a toy, a fraud, a monster, or a threat. I see a system moving from promise to structure. The promise appears to be universal

access. The structure forming underneath is cost, licensing, institutional advantage, and the metering of machine attention.

My encounter with AI began almost as a joke. I fed my poetry into a language model to see what kind of nonsense came back. Some nonsense came back. But something else came back too. The thing was oddly well read. It knew Roman Jakobson. It knew the Russian formalists, obscure literary theorists and all the rest. It recognized formal pressures, visual lineages, mythic residue, civic structure, musical pattern, and modernist inheritance. It could be wrong with great confidence, which made it resemble people I have known. But it could also apply pressure. That was the surprise.

I did not use AI because I thought it could replace the writer. I used it because I wanted to see what happened when a system trained on human language encountered a body of poems built to resist easy reading. These were not prompts. They were poems, many of them written years earlier, some published in literary journals and carrying a public record before the machine ever saw them. AI did not write them. It became a pressure chamber, proofreader, interlocutor, structural critic, and tireless counter-reader. It was useful because I could argue with it. It was useful because I could reject it. It was useful because I knew enough to know when it was smoothing the work into something false. The currently inexpensive AI ecosystem did not replace expertise. It allowed more of the working process to stay in one pair of hands.

I had edited something like thirteen books before *Between Here and the Sound of a Groaning Interstate & The Groaning Book: A Poetics*, so I knew the usual apparatus. A book is never only written. It is edited, proofread, designed, laid out, printed, distributed, corrected, carried, and eventually, with luck, read. Every book costs time. The question is whose time, and how that time is paid for. On other books, there had been designers, proofreaders, copy editors, indexers, printers, and production support. Those people did real work, and that work cost real money. Layout and design alone can become serious expenses, especially when the book has formal complications.

With this book, the whole of the apparatus moved back through me. AI did not erase the labor. It helped me carry it inside one long working process. AI did not democratize authorship in my case. It democratized access to the support structure around authorship.

The public conversation about AI keeps returning to monster language. Is the machine conscious? Can it become evil? Can users coax it into offensive combinations of words? Will it replace writers? Will it replace teachers? Will it replace thought? Will it destroy humanity? Some of those questions are serious. Some are performative. Some are ways of avoiding the more ordinary pressure: access to the powerful versions of the future.

AI is not making human moral choices. It is being built, trained, constrained, priced, marketed, and distributed by human systems. In this light, the question shifts from

what the machine can do to who will be allowed to use it deeply, at what price, and through what institutions. The danger may not be that AI becomes Frankenstein's monster. The danger may be that AI becomes a subscription wall around machine attention.

Brian Merchant's recent argument published in *Wired* about AI as the ultimate tech bubble is useful because it places the current mania inside an older pattern. Electricity, aviation, radio, and the internet all moved through speculative phases. A bubble does not prove a technology is useless. It means the story around the technology has outrun the working business model. When the speculative version breaks, the practical version remains.

If the AI bubble bursts, it will not take AI with it. The dot-com crash did not erase the web. It cleared away companies, fantasies, valuations, and bad business models. What remained became the infrastructure for much of what now surrounds us: social media, e-commerce, search, streaming, and the platform economy.

The same will happen here. The technology will move toward practical reality. The data centers, chips, energy contracts, models, interfaces, and habits of use are real. What follows will come down to who gets access to what, and who can pay for it. Most people will not keep access to the kind of deep machine attention that is now cheap, or even free. The current price is not the real price. Tech companies are still subsidizing what users ask these systems to do through capital, cloud partnerships, infrastructure buildout, energy systems, and the race to normalize dependence.

My own use depended on that gap between price and cost. AI did not make my book inexpensive to produce. I paid for it in thousands of hours. What changed was the distribution of the other costs. Work that might have gone to designers, proofreaders, editors, and production support moved back through me, with AI helping me carry more of it. The machine attention itself was subsidized by the companies behind the LLMs. Only the printing and distribution sat outside that arrangement, and those relied on an older infrastructure: e-commerce, print-on-demand, metadata systems, online retail, and the book industry the web has already remade.

The book depended on two technological moments at once: the mature web, which made printing and distribution possible through e-commerce and print-on-demand, and the immature AI economy, which made deep machine attention temporarily cheap. Part of the illusion of the present moment is that machine attention is nearly costless. The cost is hidden from the user.

Human attention costs time, discipline, memory, judgment. Machine attention costs compute, infrastructure, energy, training, capital. Social attention costs advertising money. Institutional attention costs affiliation, licenses, subscriptions, permissions. Attention always has a cost, and someone has to pay the bill.

Many technologies begin with a feeling of openness and end inside a turnstile. The early web felt free, eccentric, ugly, personal, and strange. Social media promised connection. Research databases promised scholarship. Adobe promised creative empowerment. Streaming promised access to media that once had to be bought in physical form. Search promised discovery. Each of these became something else. Access became subscription. Connection became metered reach. Knowledge became licensed database. Software became rent. Discovery became advertising. Creative work became platform dependency.

LexisNexis and Westlaw turned legal research into a tool almost nobody outside a firm, university, or library can afford to use seriously. JSTOR and ProQuest turned scholarship into a system mediated by institutional affiliation. Adobe turned the standard tools of professional design into a monthly bill. Calling it corporate villainy is too easy. It is what happens when powerful tools survive their speculative phase and become transactional: priced, licensed, rented, bundled, and sold back to the people who have learned to need them.

A database gates information. AI gates the ability to work across information. The inequality is not only who can read the archive. It is who gets a machine licensed and trained to move through the archive with them.

Social media gave the clearest preview. I have been paying to place poems from the book in front of readers on Facebook, essentially serializing the work in the oldest sense. The mechanics tell the story. I post a poem. People respond. Some share. A small percentage of my followers see the post. A smaller percentage of a sharer's followers see the share. Each act of human circulation enters another algorithmic bottleneck. The system lets enough through to prove the audience is there, then sells reach back to the writer. Social media allows connection, then taxes its reach. Social media sells access to human attention. AI will sell access to machine attention.

Large language models create a space in which a single user can manufacture deep attention. The models can read, summarize, compare, revise, proof, search, analyze, pressure-test, translate, and return again and again to the same project. That is machine attention, and deep machine attention is expensive.

The future will probably not divide into people who have AI and people who do not. Almost everyone will have some AI. The divide will be in depth: context, memory, archives, privacy, reliability, and the ability to stay with difficult work. Institutions will have licensed datasets, long context windows, private models, citation systems, archival access, and technical support. Independent users will have caps, ads, weaker models, generic answers, and subscription creep, or they will pay a premium for the luxury of institutional access.

Government may complicate this picture. If AI becomes central to education, search, law, public administration, medicine, and cultural memory, public investment is plausible. So are regulated utilities, library and university partnerships, procurement-driven access, and licensing regimes. A gate remains. What changes is who builds it, who pays for it, who stands inside it, and who is left asking for entry.

This is where the moral weather around AI starts to look strange. Many writers, artists, and academics are reacting to AI with revulsion. Some of that revulsion is earned. The training-data questions are real. The replacement rhetoric is real. The corporate arrogance is real. Refusal may be ethically coherent. It may also be strategically costly.

The cheap window for unaffiliated experimentation is closing while the revulsion is still being articulated. Many of the people most capable of using the tool seriously are refusing to touch it while it remains broadly available. Later, those same people may receive licensed, archive-connected, professional-grade AI through universities, publishers, media companies, law firms, and research centers. Independent writers will be offered the consumer versions.

Moral outrage is not the same thing as ethics. Ethics has memory. Outrage often moves on. That does not make the objections to AI false. It means the market will not wait for disgust to resolve itself into policy. While the shuddering continues, the access model is hardening.

In 2000, Lars Ulrich became the public villain of the Napster fight. The internet demonized him for arguing that Metallica's music should not be freely copied just because copying had become easy. He was rich. He was famous. He was an easy target. But the structure of his argument was simple. Technology makes copying easy. That does not make it ethical.

Now many writers, artists, journalists, and academics are making a version of the same argument about AI training. They are not necessarily wrong. Lars, the old villain, has become the new argument. The leaked Metallica song at the center of the Napster controversy was called "I Disappear." The irony is almost too exact.

A quarter century later, disappearance may turn out to be the second fear of writers in the age of AI. The first fear is that the work will be taken. The second fear, harder to see right now, is that the work will be left outside the machine-readable archive. Today's scandal is unauthorized inclusion. Tomorrow's scandal may be exclusion without consequence.

If AI becomes part of search, recommendation, citation, scholarship, teaching, and cultural memory, then absence from machine-readable systems becomes a form of disappearance. The work may exist. It may sit on a shelf. It may be available through a publisher. It may carry an ISBN. But if future systems cannot find it, summarize it,

recommend it, cite it, or place it in context, it becomes culturally uncalled and invisible. The first fight was against being scraped. The second fight may be against being forgotten.

One counter-move remains available now, before the door narrows: open placement. A Creative Commons release on a public archive makes the work available as a public record. The book remains a book. It also becomes searchable, citable, findable, and machine-readable. This is not surrender to AI. It is a wager on availability.

If the future archive is partly machine-mediated, then leaving the work where human and machine attention may both find it is a form of survival.

Between Here and the Sound of a Groaning Interstate & The Groaning Book: A Poetics was made during a brief period when powerful AI systems were available to an unaffiliated writer at consumer cost. The poems had already existed before the machine. Some had appeared in literary journals. The AI encounter did not originate the work. It altered the field of attention around it.

This was not because the machine made expertise unnecessary. It was because an unusual human formation met the tool at the right time. Poet. Editor. Bookmaker. Designer. Civic administrator. Early-web adapter. Student of poetics. Executive of public institutions. Humane officer. Fundraiser. Museum and library worker. Someone who had spent years close to artists, soldiers, donors, boards, trustees, courts, animals, books, and public language. AI entered that existing concentration. It did not create the lens. It gave the lens a surface.

The book is not evidence that AI can replace a writer. It is evidence that, for a brief interval, a writer outside the institution could borrow enough machine attention to build the apparatus around a difficult work. That interval may be short.

The same tool, refined and folded into professional workflows, will be used by publishers, agents, marketing departments, academic presses, metadata systems, and review institutions. Authors may be warned against the machine while the industry uses the machine around them. The tools that briefly allowed an outsider to assemble parts of an institution around a book will be folded back into institutions.

The deep version of AI will probably not disappear. It will become ordinary and expensive. The closure will not announce itself as closure. It will arrive as a subscription, a license, an enterprise tier, a university agreement, a government contract, a database integration, a professional workflow, a feature inside another product. By then, the democratic moment will already have passed. What remains is access, priced accordingly.

The next paywall will think.

About the Author: KC Clarke is the author of *Between Here and the Sound of a Groaning Interstate* and *The Groaning Book: A Poetics*, and of *Wolves and Flax: The Prior Family in the Cuyahoga Valley Wilderness*. His professional life has included leadership in literary centers, museums, libraries, and nonprofit institutions, as well as editorial and creative direction for numerous books. He writes from within the civic and cultural infrastructure this essay describes.